WATER RESOURCES RESEARCH, VOL. 37, NO. 9, PAGES 2353–2364, SEPTEMBER 2001

Downscaling of radio brightness measurements for soil moisture estimation: A four-dimensional variational data assimilation approach

Rolf H. Reichle,[1] Dara Entekhabi, and Dennis B. McLaughlin

Ralph M. Parsons Laboratory, Department of Civil and Environmental Engineering
Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Abstract. This paper investigates the feasibility of estimating large-scale soil moisture profiles and related land surface variables from 1.4 GHz (L-band) passive microwave measurements, using variational data assimilation. Our four-dimensional assimilation algorithm takes into account both model and measurement uncertainties and provides dynamically consistent interpolation and extrapolation of remote sensing data over space and time. The land surface hydrologic model which forms the heart of the variational algorithm was expressly designed for data assimilation purposes. This model captures key physical processes while remaining computationally efficient. We test our algorithm with a series of synthetic experiments based on the Southern Great Plains 1997 Hydrology Experiment. These experiments provide insights about three issues that are crucial to the design of an operational soil moisture assimilation system. Our first synthetic experiment shows that soil moisture can be satisfactorily estimated at scales finer than the resolution of the brightness images. This downscaling experiment indicates that brightness images with a resolution of tens of kilometers can yield soil moisture profile estimates on a scale of a few kilometers, provided that micrometeorological, soil texture, and land cover inputs are available at the finer scale. In our second synthetic experiment we show that adequate soil moisture estimates can be obtained even if quantitative precipitation data are not available. Model error terms estimated from radio brightness measurements are able to account in an aggregate way for the effects of precipitation events. In our third experiment we show that reductions in estimation performance resulting from a decrease in the length of the assimilation time interval are offset by a substantial improvement in computational efficiency.

1. Introduction

Soil moisture plays a major role in the global hydrologic cycle, principally through its effect on the partitioning of energy and precipitation at the land surface. As a result, soil moisture is a key variable for weather and climate prediction, flood forecasting, and the determination of groundwater recharge. It is well known that soil moisture is difficult to measure over the scales needed for these applications. In situ soil moisture profiles measured at point locations vary greatly and are inadequate for characterizing regional variations. Remotely sensed radio brightness measurements (e.g., 1.4 GHz L-band passive microwaves) provide better spatial coverage. However, they are sensitive only to soil moisture in the top 5 cm of the surface layer, and they do not provide direct information about the soil moisture profile.

The limitations of current techniques for measuring soil moisture can be partially overcome if data assimilation techniques are used to combine radio brightness measurements with other relevant information, such as micrometeorological measurements, land cover data, and soil textural maps. These complementary sources of information can be related with physically based models that describe the connection between land surface states and measured radio brightness. Since the measurements may be noisy and the models are only approximate, the data assimilation procedure must explicitly account for both measurement and model errors. This prevents the procedure from relying excessively on either source of information.

The models used in this study describe the state of the land surface in terms of the near-surface soil moisture and soil temperature and the temperature in the overlying vegetation canopy. The moisture and temperature states are constrained by mass and energy balances which are formulated over an extensive grid of vertical soil columns. A simple radiative transfer model relates the measured radio brightness signal from a given area to the model states in the associated columns. The soil and canopy states are adjusted to provide a best fit to the measured radio brightness, with consideration given to the possible impact of measurement and model errors. This approach provides a true four-dimensional soil moisture data assimilation capability.

Data assimilation is distinguished from more traditional hydrologic model calibration primarily by its objectives. In data assimilation the emphasis is on the estimation of hydrologic states from measurements. The model is viewed as a tool in this process, rather than as the primary object of interest. As a result, model selection should consider a number of factors,

[1]Now at NASA Goddard Space Flight Center, Greenbelt, Maryland, USA.

Paper number 2001WR000475.
0043-1397/01/2001WR000475$09.00

including computational efficiency and robustness as well as physical realism.

Given certain assumptions, the assimilation process can be reduced to the solution of a constrained least squares estimation problem. This problem is difficult because the number of unknowns can be very large and the constraining model is highly nonlinear. Many of the existing soil moisture assimilation studies focus on one-dimensional problems [*Katul et al*., 1993; *Parlange et al*., 1993; *Entekhabi et al*., 1994; *Calvet et al*., 1998; *Galantowicz et al*., 1999; *Castelli et al*., 1999]. Since in one dimension the number of nodes is usually very small and the computational resources are not a limiting factor, these studies apply optimal assimilation techniques such as the Kalman filter [*Gelb*, 1974]. In most cases, these optimal algorithms are able to quantify the accuracy of their estimates by explicitly computing error bars, or more generally, the estimation error covariance matrix, which typically varies over both time and space.

A second category of soil moisture assimilation studies confronts the problem of estimating horizontal as well as vertical variations in soil moisture. In order to deal with computational limitations these studies use simplified estimation algorithms. *Houser et al*. [1998] discuss various algorithms (including Nudging and Statistical Interpolation). However, these methods are suboptimal because they ignore the spatial and temporal evolution of the estimation error covariance matrices and use static covariances instead.

Variational data assimilation methods used in meteorological and oceanographic applications [*Thépaut and Courtier*, 1991; *Daley*, 1991; *Bennett*, 1992; *Courtier et al*., 1993; *McLaughlin*, 1995] offer the possibility of achieving the optimal performance of Kalman filters with the computational efficiency of suboptimal methods. This is possible because variational methods do not explicitly evaluate the large error covariance matrices which are propagated by Kalman filters. Instead, variational algorithms simultaneously process all data within a given time period (or "assimilation interval") and implicitly take dynamic error information into account by propagating an adjoint variable (section 2). Various computational enhancements developed in recent years have greatly improved the practical potential of the variational approach [*Bennett*, 1999].

This paper considers the feasibility of applying variational methods to large-scale land assimilation problems. It also presents some important results relating to the design and the implementation of operational assimilation algorithms. We begin in sections 2 and 3 with brief reviews of the assimilation method and the land model. Details are provided in the cited references, including *Reichle* [2000]. In section 4 we describe several synthetic experiments that we later use to investigate design issues. Section 5.1 demonstrates that coarse-scale brightness data are useful for the estimation of soil moisture at finer scales (downscaling). Section 5.2 shows that reasonable soil moisture estimates can be obtained from L-band passive microwave observations, even if accurate quantitative precipitation data are not available. Section 5.3 investigates the length of the assimilation interval and related initialization issues. In section 6, we summarize our results and their limitations.

2. Variational Data Assimilation

We begin by recognizing that model predictions and measurements both provide useful information about the actual state of the soil. The primary task of data assimilation is to combine these two sources of information in an optimum way. The models used for data assimilation are typically based on mass, momentum, or energy balance equations which describe the temporal and spatial evolution of specified state variables. In the application of interest here the primary states are soil moisture, soil temperature, and canopy temperature. These variables may be approximated by a set of spatially discretized states defined at the cells of a three-dimensional computational grid (Figure 1).

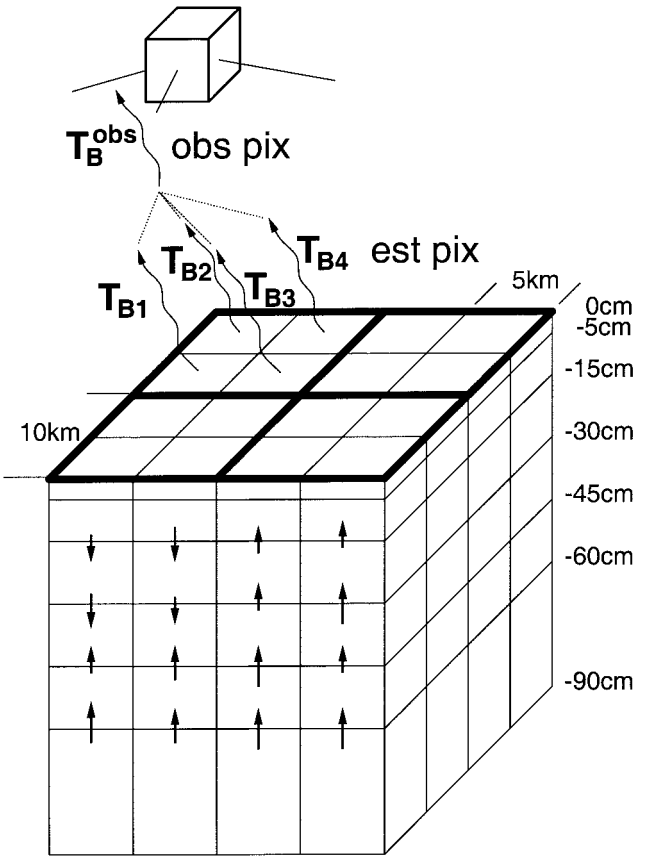

Figure 1. Schematic of the model grid with typical length scales. Soil moisture estimates are obtained on the 5-km scale of the estimation pixels (fine outline). For this example remotely sensed brightness measurements are available on the 10-km scale of the observation pixels (thick outline).

If the spatially discretized states are assembled in a time-dependent N_Y dimensional vector $Y(t)$ the hydrologic model can be expressed as

$$\frac{\partial Y}{\partial t} = \varphi(Y) + \omega, \qquad Y|_{t=0} = Y_0(\beta). \tag{1}$$

We account for model errors by treating the energy and moisture forcing terms in the model as temporally and spatially correlated random fields. These model errors are assembled in the N_Y dimensional vector ω. The nonlinear operator φ includes all deterministic forcings such as observed micrometeorologic inputs. The initial condition depends on a random vector β, which may have a dimension less than N_Y. The random variables which appear in (1) are defined in terms of their first and second-order moments (means and covariances). In particular, the specified prior (or first guess) mean values $\bar{\omega}$

and $\bar{\beta}$ and prior (or first guess) covariances C_ω and C_β summarize the available information about the random errors prior to assimilating the radio brightness measurements.

The radio brightness measurements to be assimilated by the estimation algorithm are indirectly related to the moisture and temperature states in the hydrologic model. It is convenient to assemble all brightness temperature measurements available over a specified time period $[0,\ t_f]$ in an N_Z dimensional vector Z. We call this time period the "assimilation interval." The kth element (Z_k) of Z is the scalar measurement obtained at time $t_k \in [0,\ t_f]$ over an area centered on location x_k. With these definitions we can construct a measurement equation which relates the measurement and state vectors:

$$Z = M[Y] + \upsilon. \tag{2}$$

The measurement operator $M[Y]$ is an N_Z dimensional vector. Element k of this vector describes the relation between a particular measurement Z_k and the state vector $Y(t_k)$ defined at measurement time t_k. In our land surface application the elements of the measurement operator $M[Y]$ are derived from a nonlinear radiative transfer model (section 3.1). The N_Z dimensional vector υ in (2) accounts for additive measurement errors and is defined in terms of its first and second moments. Here we assume that υ is zero mean with a specified covariance matrix C_υ.

The data assimilation algorithm produces a state estimate which balances the effects of uncertain model errors, initial conditions, and measurement errors. This can be achieved if the uncertain inputs included in the problem formulation (the initial states and the model errors) are adjusted to minimize a weighted least squares (Bayesian) performance function J.

$$J = (Z - M[Y])^T C_\upsilon^{-1}(Z - M[Y]) + (\beta - \bar{\beta})^T C_\beta^{-1}(\beta - \bar{\beta}) + \int_0^{t_f}\int_0^{t_f} \omega(t')^T C_\omega^{-1}(t', t'')\omega(t'')\, dt'\, dt'' + 2\int_0^{t_f} \lambda^T\left(\frac{\partial Y}{\partial t} - \varphi(Y) - \omega\right) dt. \tag{3}$$

Throughout the text, superscript T denotes the vector or matrix transpose. The performance function minimizes the aggregate error over the assimilation interval. The first term in this function takes into account the misfit between the data vector Z and the measurement predictions $M[Y]$, normalized by the measurement error covariance C_υ (which has dimension $N_Z \times N_Z$). The second and third terms penalize normalized deviations of the uncertain inputs from the specified prior mean values. For convenience we assume that the mean model error $\bar{\omega}$ is zero. This assumption could be relaxed if available information suggested the presence of a systematic bias (non-zero mean) in the model error. The final term in (3) is obtained by adjoining the state equation to the performance index with the time-dependent N_Y dimensional Lagrange multiplier vector λ. The model constraint ensures that the state estimates produced by the data assimilation procedure will be dynamically consistent.

The optimal estimates of the uncertain inputs and states are obtained by setting the first variation of the adjoined performance function (3) equal to zero. This yields a set of so-called Euler-Lagrange equations, which constitute a two-point boundary value problem [*Reichle et al.*, 2001].

$$\frac{\partial \hat{Y}}{\partial t} = \varphi(\hat{Y}) + \hat{\omega} \qquad \hat{Y}|_{t=0} = Y_0(\hat{\beta}), \tag{4}$$

$$-\frac{\partial \lambda}{\partial t} = \frac{\partial \varphi}{\partial Y}\bigg|_{\hat{Y}}^T \lambda + \frac{\partial M}{\partial Y}\bigg|_{\hat{Y}}^T [\delta] C_\upsilon^{-1}(Z - M[\hat{Y}]) \qquad \lambda|_{t=t_f} = 0, \tag{5}$$

$$\hat{\omega} = \int_0^{t_f} C_\omega(t, t')\lambda(t')\, dt' \qquad \hat{\beta} = \bar{\beta} + C_\beta \frac{\partial Y_0}{\partial \beta}\bigg|_{\hat{\beta}}^T \lambda|_{t=0}. \tag{6}$$

In these equations, estimates are denoted with a hat and $[\delta] \equiv \mathrm{diag}\ (\delta_1,\ \delta_2,\ \ldots,\ \delta_{N_Z})$ is a diagonal matrix with the kth diagonal element $\delta_k \equiv \delta(t - t_k)$, where $\delta(t - t_k)$ is the scalar Dirac delta function for time t_k. Note that $\partial\varphi/\partial Y$ and $\partial M/\partial Y$ are $N_Y \times N_Y$ and $N_Z \times N_Y$ matrices, respectively.

The first, or forward, Euler-Lagrange equation (4) expresses the fact that the estimates obey the state equation (1). The second, or adjoint, equation describes how measurement information obtained at a given time propagates backward to earlier times. Note that this equation is forced by differences between measurements and the corresponding model predictions. The third, or update, set of equations (6) relates the estimates of the uncertain inputs to the adjoint variable. This set of equations indicates that the fundamental unknowns that must be estimated are the elements of the model error $\hat{\omega}$ and the initial condition parameter vector $\hat{\beta}$. Once these unknowns are estimated, the forward equation may be used to derive the state estimate $\hat{Y}$ at any time or location within the assimilation interval. Note that the adjoint variable λ contains all of the measurement information needed to derive the estimates $\hat{\omega}$ and $\hat{\beta}$. These estimates and the state estimates that depend on them can be obtained without explicit computation of the state estimation error covariance matrix.

The coupled set of nonlinear Euler-Lagrange equations provided above must be solved with an iterative numerical algorithm. Although a number of possibilities are available, we have found that the best combination of efficiency and performance is provided by the iterated indirect representer algorithm [*Bennett et al.*, 1998; *Reichle et al.*, 2001]. This algorithm relies on a series expansion of the estimated fields. The basis functions of the expansion are the linearized prior cross covariances between the measurement predictions and the states. These basis functions are derived from the Euler-Lagrange equations rather than specified a priori.

Since the iterated indirect representer algorithm propagates covariance information implicitly (through the Euler-Lagrange equations) rather than explicitly (with a large covariance matrix) it is very efficient. In particular, this algorithm is able to provide statistically optimal estimates without the simplifications that have been used in other large-scale soil moisture estimation applications.

3. A Land Surface Hydrologic Model for Data Assimilation

A hydrologic model for land surface data assimilation must capture key physical processes while remaining efficient enough to make large-scale optimal estimation computationally feasible. These are potentially conflicting requirements that need to be traded off when a model is selected. On the

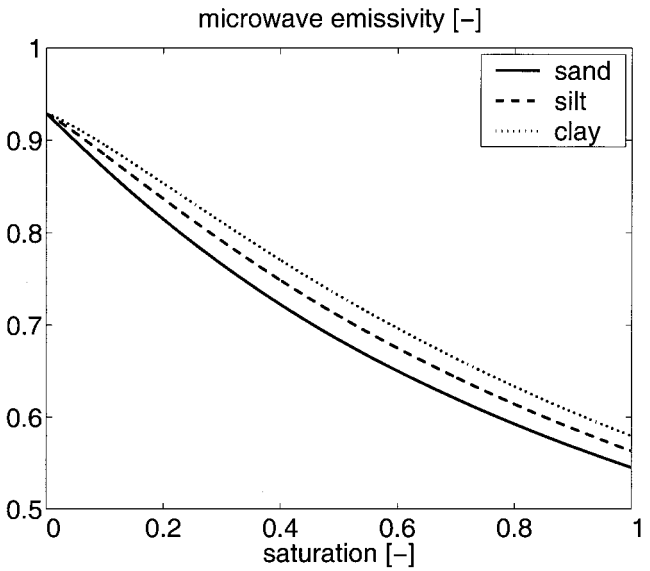

Figure 2. Bare soil microwave emissivity versus saturation [*Ulaby et al*., 1986].

basis of the variety of models described in the literature, we have developed a land surface scheme especially designed for data assimilation purposes. This model is described in detail by *Reichle* [2000]. Its most important features are summarized in the following sections.

3.1. Vertical Moisture and Energy Transport

Our model of coupled moisture and heat transport is a typical soil-vegetation-atmosphere transfer scheme (SVAT). Vertical soil moisture transport is governed by a spatially discretized one-dimensional (vertical) version of Richards' equation, with soil hydraulic properties described by relationships from *Clapp and Hornberger* [1978]. Figure 1 shows the six subsurface layers used in the Richards' equation discretization. These are located at 0–5 cm, 5–15 cm, 15–30 cm, 30–45 cm, 45–60 cm, and 60–90 cm. The lower boundary condition for the soil moisture equation is defined by gravity drainage.

In order to achieve maximum computational efficiency we describe temperature dynamics with a one-layer force-restore method [*Hu and Islam*, 1995] rather than the full heat equation. This approach, which focuses on diurnal fluctuations in the upper soil layer, is ideal for our remote sensing application. The deep soil temperature in the force-restore model is specified as the seasonal average of the observed air temperature.

We describe moisture and heat fluxes through the vegetation layer with a resistance network approach adapted from the Simplified Biosphere Model (SSiB) [*Xue et al*., 1991]. At the land surface we assume near-neutral atmospheric conditions. The brightness temperature is related to the land surface states with a grey body radiative transfer model [*Ulaby et al*., 1986; *Galantowicz et al*., 1999]. Figure 2 shows the bare soil microwave emissivity (which mostly determines the brightness temperature) versus soil moisture. For bare soil, there is a strong sensitivity of the microwave emissivity to soil moisture. Vegetation and soil roughness effects decrease this sensitivity, but for grasslands and crops L-band brightness data are still informative about soil moisture [*Jackson et al*., 1999].

Our land surface model has been successfully tested by comparing its predictions with measurements from the BARC data set [*Jackson et al*., 1997; *Reichle*, 2000]. On the basis of these tests we believe that the model described here is sufficiently accurate and computationally efficient to form the basis for an operational soil moisture data assimilation algorithm.

3.2. Horizontal Variability and Downscaling

Although the model used in our data assimilation procedure provides for spatial variability over all three dimensions, the fluxes and length scales in the vertical and horizontal dimensions are much different. This feature can be exploited to improve computational efficiency without any significant sacrifice in accuracy. In particular, we assume that lateral (horizontal) moisture and heat fluxes in the unsaturated zone are negligible. As a result, horizontal structure in our soil moisture estimates reflects spatial correlation in micrometeorological inputs, land cover, and soil texture rather than horizontal transport. This assumption is reasonable for terrain with moderate relief over the spatial scales under consideration here. It enables us to break the model domain into a grid of one-dimensional vertical cells or "estimation pixels," as illustrated in Figure 1. We seek estimates of the soil moisture in the seven nodes associated with each pixel and estimates of the canopy temperature and the soil temperature in the surface layer of each pixel.

Satellite observations of L-band brightness temperatures will likely be available only over relatively large spatial scales. For present purposes we define each pixel resolved in the brightness image to be an "observation pixel." Over the United States the micrometeorological data and other inputs defined over the estimation pixels of our model grid are typically available on scales of 1 km or less. By contrast, the observation pixel scale for a space-borne passive L-band sensor will be of the order of 50 km in the near future and possibly 10 km in 10 years. Clearly, it would be desirable to derive soil moisture estimates at the finer scale of the estimation pixels, even when brightness images are only available at the observation scale. This process is sometimes called "downscaling" to reflect the transfer of information from larger to smaller spatial scales.

The problem formulation of section 2 provides a convenient way to account for scale differences between the estimation and observation pixels. The kth element of the measurement operator $M[Y]$ describes the relationship between the brightness measurement Z_k in the observation pixel centered on x_k and the states of estimation pixels nested within this observation pixel. At the frequencies of interest here the brightness temperature in an observation pixel can be adequately approximated as the arithmetic average of the brightness temperatures in the nested estimation pixels [*Drusch et al*., 1999; *Liou et al*., 1998]. The radiative transfer model specifies how the estimation pixel brightness temperatures are related to the land surface states. When these relationships are used to define $M[Y]$, the data assimilation algorithm can estimate states at the estimation pixel scale directly from measurements at the observation pixel scale. This downscaling process is described for a simple example in Appendix A.

4. Assessing the Performance of the Algorithm With Synthetic Experiments

In section 5 we investigate a series of important design issues with data assimilation experiments based on synthetically generated radio brightness data. Synthetic experiments are ideally suited for algorithm performance tests since all of the uncer-

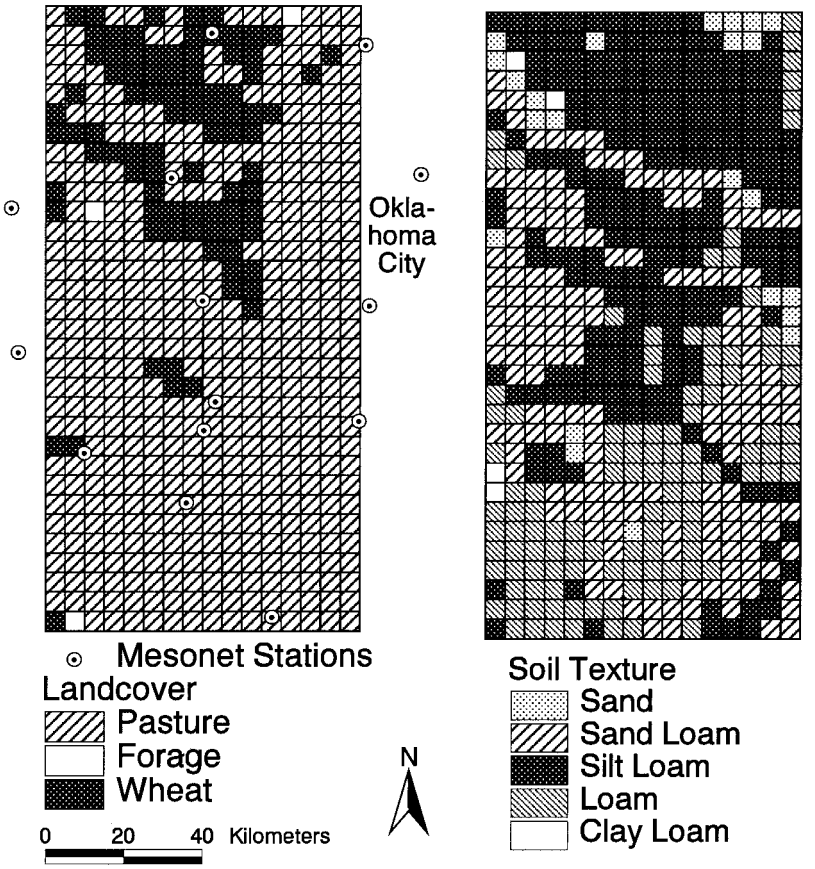

Figure 3. Area for the synthetic experiment. The horizontal grid of 16 by 32 pixels (80 km by 160 km) is shown together with the locations of the Oklahoma Mesonet stations and the (left) land cover classes and with (right) the soil texture classes.

tain inputs are known by design. There is no ambiguity about the "right answer" to be achieved. When dealing with complex data processing algorithms such as the one described here synthetic experiments are an indispensable first step toward a field application. However, it is obvious that such experiments cannot replace operational tests based on real field observations. In the near future it should be possible to carry out operational tests of soil moisture data assimilation algorithms such as the one described here. One promising option is a test using passive microwave data from the Electronically Scanned and Thinned Array Radiometer (ESTAR) deployed during the SGP97 field campaign [*Jackson et al*., 1999].

4.1. Experiment Area and Model Inputs

The synthetic test problems considered in this paper are based on the recent Southern Great Plains (SGP97) hydrology experiment in central Oklahoma [*Jackson et al*., 1999]. The 80 km by 160 km area chosen for the synthetic experiment is shown in Figure 3. We divide this area into 16 by 32 estimation pixels of 5 km by 5 km. The soil temperature and the soil moisture profile are estimated in each of these 512 pixels. Figure 3 also displays land cover data, which were obtained from the SGP97 Data Archive, and soil texture data, which were compiled from the ESSC database at Pennsylvania State University. The micrometeorologic inputs are taken from the Oklahoma Mesonet database and interpolated to the model grid with inverse-square-distance weights.

Our synthetic experiment covers a 2-week period from June 18, 1997 (day of year 169), to July 2, 1997 (day of year 183). Soil moisture estimates are derived at a basic time step of 15 min. Figure 4 shows a time series of the observed area average precipitation, with radio brightness measurement times indicated by small open circles. Note that observations are occasionally taken during or shortly after a rain event. While it is raining, the water film covering the vegetation and the soil makes it difficult to take accurate measurements of the passive L-band microwave radiation. We do not have to account for this effect in our synthetic experiments. In an operational setting, a quality control routine would be needed to screen out data affected by water films.

4.2. Synthetic Experiments

For each of our synthetic experiments we generate time-invariant initial condition parameters β and time-dependent random model error fields ω. The corresponding solution to the forward equation (1) is defined to be the set of "true" system states. These are computed every 15 min at the 5 km estimation pixel scale. The "prior" state is the solution of (1) obtained when the initial conditions and model errors are set to their prior values $\bar{\beta}$ and $\bar{\omega} \equiv 0$, respectively. This prior solution can be viewed as the "best guess" of the true states available without the benefit of radio brightness measurements.

The spatial and temporal correlation functions of the uncertain inputs are unknown a priori and very difficult to characterize. Their accurate determination for a given model and

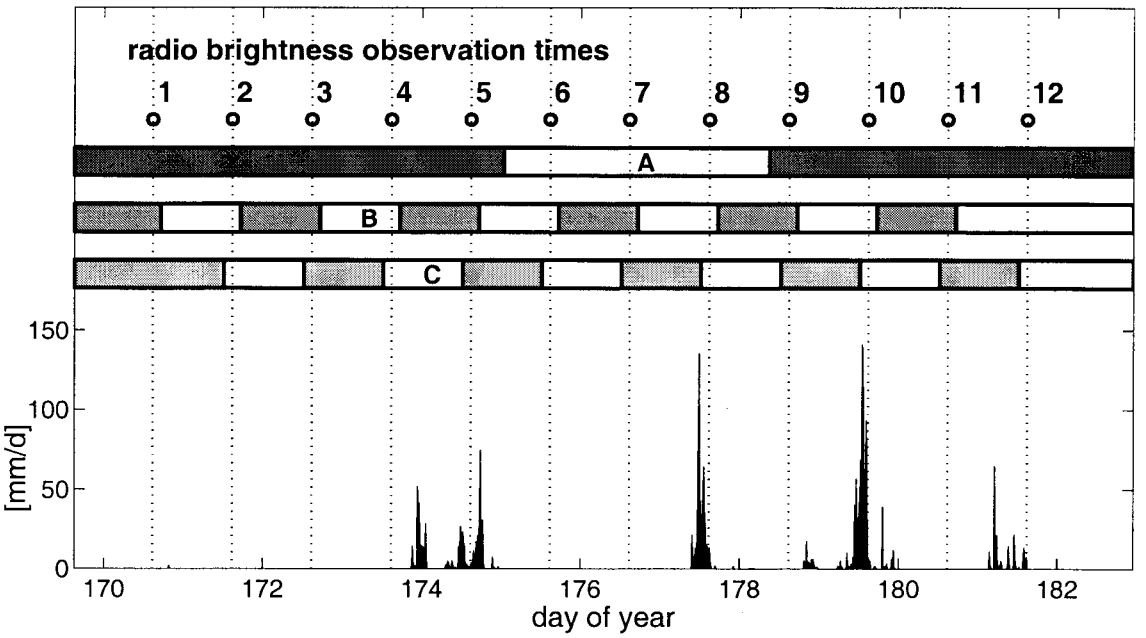

Figure 4. Areally averaged observed precipitation shown with the temporal setup of the synthetic experiments. For the downscaling experiments (section 5.1) and the precipitation withholding experiment (section 5.2) there is just one assimilation interval which covers the entire 2-week period. In experiment A of section 5.3, there are three assimilation intervals indicated on the horizontal bar labeled A. In experiments B and C, there are 12 assimilation intervals with observation times at the end and the beginning of the intervals, respectively. These are indicated on the horizontal bars labeled B and C.

field setting constitutes a research project in its own right and is well beyond the scope of this paper. Here, we only aim to prove the concept of soil moisture assimilation and downscaling with the synthetic experiments. This does not critically depend on the exact shapes and scales of the correlation functions, and we specify conditions that in our experience are appropriate for the experiment area and our model.

The initial conditions and model errors generated in our synthetic experiments are normally distributed random fields with exponential spatial and temporal correlation functions. To make the estimation more robust, we prescribe the shape of the soil moisture profile at the initial time, but the total amount of water that is stored across the column is uncertain. The shape of the profile can for example be the best estimate at the final time of a previous assimilation interval, a profile that results from a spin-up integration, or simply a hydrostatic profile. In our experiments the initial condition of the top node soil saturation has a mean value of 0.5, an error standard deviation of 0.07, and a correlation length of 50 km. After the initial time, changes in the surface fluxes cause the moisture profile to deviate from its initial shape. The initial upper layer soil temperature is set equal to the initial air temperature and is assumed to be known perfectly (the upper layer soil temperature memory is only a few hours and has little impact on longer-term estimates).

Model errors are represented as unknown fluxes in the near-surface soil moisture, soil energy, and canopy energy balance equations. We assume that each of these errors is zero mean with a standard deviation of 50 W m^{-2}. The model error correlation lengths are all 6 km and their correlation times are all 10 hours.

The true brightness temperatures are obtained by running the true states through the radiative transfer model at the chosen observation pixel scale. The vector Z of synthetic brightness measurements is obtained by adding random measurement errors. The assimilation algorithm estimates the land surface states from the noisy data Z and the prior solution. The algorithm's performance may be measured in terms of the difference between the true and estimated states at a given time and location (the estimation error). In some cases, it is convenient to consider the root-mean-square (rms) error averaged over all pixels in the study area. The algorithm's estimation error should be compared to the prior error, which is the difference between the true and prior states. If the measurements are informative the rms estimation error should be smaller than the rms prior error.

The synthetic brightness measurements used in our experiments are generated at several different scales, corresponding to different downscaling scenarios. In each case a daily synthetic brightness value is generated at every observation pixel in the model domain at 0945 LT. The random measurement errors added to the brightness temperature values are spatially and temporally uncorrelated with a standard deviation of 5 K. This observation time and level of uncertainty are typical of the SGP97 field experiment. The absence of spatial correlation is not a constraint imposed by the algorithm but is a simplification adopted for convenience.

Figure 4 summarizes the temporal setup of the experiments considered in this paper. For the "reference experiment" the observation pixel is the same size as the estimation pixel (5 km on a side), precipitation measurements are provided to the assimilation algorithm, and there is only one assimilation interval spanning the entire two week period for which data are available. Results from the reference experiment are described in detail by *Reichle* [2000]. In this paper we focus on modifi-

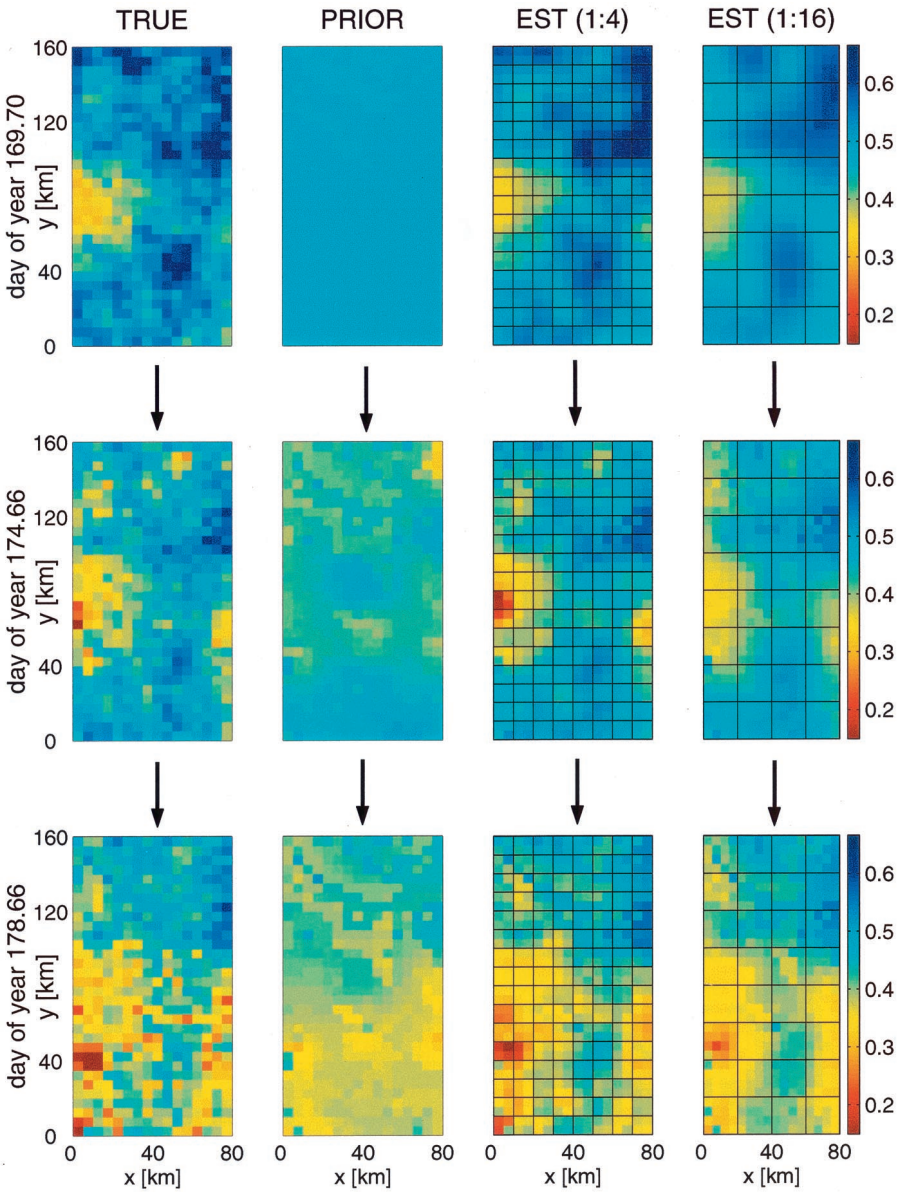

Plate 1. True and prior top node saturations (in the first two columns) at three different times (in the three rows). Estimated top node saturations for downscaling ratios of (1:4) and (1:16) are also shown (in the third and fourth columns). The observation pixels used to compute the downscaled estimates are indicated with solid black grid lines.

cations to the reference experiment that illuminate some important design issues.

5. Results and Discussion

In the following subsections we investigate three design issues which are relevant to the task of estimating soil moisture from remotely sensed radio brightness measurements. The first of these concerns the feasibility of downscaling, defined as the estimation of soil moisture and other land surface states at spatial scales finer than the scale of radio brightness observations. The second issue concerns the feasibility of estimating soil moisture from radio brightness measurements in the absence of precipitation measurements. The third issue considers the effect of the assimilation interval length on the accuracy and computational demands of the assimilation algorithm.

5.1. Downscaling

As outlined in section 3 and Appendix A, we can effectively increase the resolution of radio brightness images by making use of the fact that micrometeorological, land cover, and soil texture information are typically available at finer scales than the brightness data. The spatial structure conveyed by these finer-scale model inputs provides the information needed to resolve soil moisture variations over scales smaller than the observation pixel scale. We can investigate this process in detail by considering estimates obtained from a series of synthetic experiments with progressively coarser observation pixels. Each of these is the same as the reference experiment discussed in section 4 except that the size of the observation pixel is changed.

Recall that the estimation pixel used in all of our synthetic

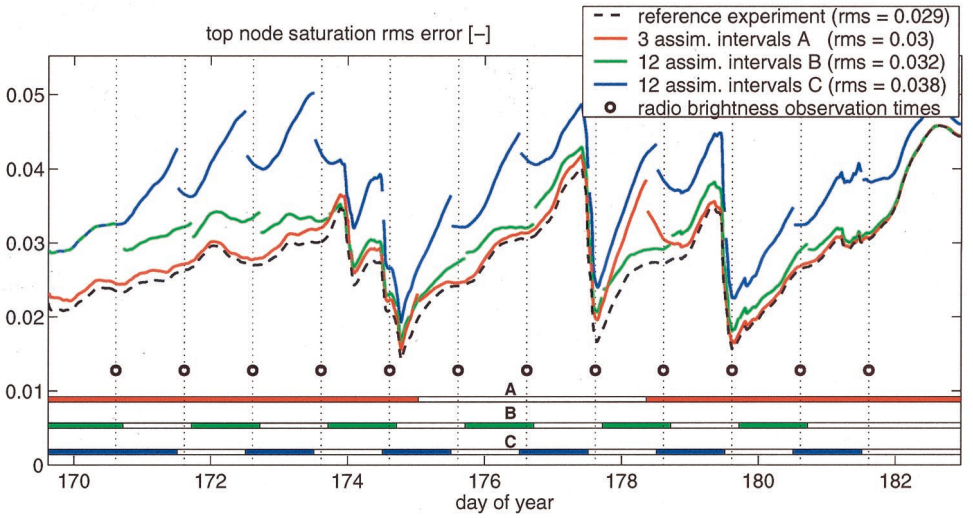

Plate 2. Areally averaged root-mean-square (rms) errors for multiple assimilation intervals. The rms errors in the estimated top node saturation are shown for the single interval reference experiment, one experiment (A) with three assimilation intervals, one experiment (B) with twelve assimilation intervals with measurements located near the end of each interval, and one experiment (C) with twelve assimilation intervals with measurements located near the beginning of each interval. Discontinuities in estimates occur at interval boundaries because the initial condition estimated from radio brightness measurements in interval n is generally different than the final value estimated from measurements in interval $n - 1$. Sharp decreases in rms error occur when measurements follow soon after precipitation events (Figure 4). The legend provides the temporal mean of the rms errors.

experiments has a length of 5 km (Figure 1). For our downscaling experiments we generate synthetic brightness measurements over observation pixels with lengths of 10 and 20 km, respectively. In the first case, each observation pixel contains four estimation pixels, while in the second case each observation pixel contains 16 estimation pixels. We refer to these two alternatives as the (1:4) and (1:16) downscaling scenarios, respectively.

In Plate 1 we compare the true and prior top node saturations (shown in the first and second columns, respectively) to the top node saturation estimates obtained from each of the downscaling scenarios (shown in the third and fourth columns). Recall that the prior solution corresponds to a simulation without updates from the brightness observations. Each row in Plate 1 corresponds to a different time. Note that soil moisture estimates are also obtained at the six nodes below the surface layer, providing a complete soil moisture profile at each pixel. The general patterns at depth are similar to those observed in the surface layer.

Plate 1 demonstrates that the algorithm can adequately estimate the large-scale spatial distribution of the saturation in both of the downscaling scenarios. Structures at scales well below the scale of the observations can be resolved satisfactorily. This means that radio brightness images with resolutions of a few tens of kilometers are useful, even if the estimation scale of interest is of the order of a few kilometers, provided that fine-scale information is available on the micrometeorologic forcings, land cover, and soil texture.

5.2. Assimilation Without Precipitation Data

Of all model inputs, precipitation is the one which most dominates soil moisture. At the same time, precipitation is also the input with the highest uncertainty. Rain gauge observations (rain depths over specified time intervals) are point measurements, and interpolation to larger areas is difficult and error prone. Reflectance measurements from radar sensors provide greater coverage than point rainfall data but they are only indirectly related to precipitation. Since precipitation measurements are so important in soil moisture assimilation, it is useful to take a closer look at their role in the assimilation process.

One way to do this is to examine the effect of withholding precipitation data. We do this by repeating the reference experiment with all precipitation inputs set equal to zero while assimilating the daily synthetic brightness observations that have been generated at 0945 LT at each observation pixel. In this case the effects of precipitation can be accounted for by an increase in the estimated model error for the soil moisture state. In order to enable the model error to increase sufficiently during rain events, we increase the model error standard deviation for the soil moisture flux to 2.7 mm h^{-1} at all pixels when precipitation occurs somewhere in the domain. That is, we tell the algorithm when it might be raining (a binary indicator), but we do not specify the amount or location of rain falling.

Figure 5 shows the root-mean-square (rms) top node saturation error when quantitative precipitation data are withheld from the assimilation algorithm. For comparison, we also show the rms estimation error for the reference experiment, which includes precipitation data. Figure 5 indicates that even when all quantitative precipitation information is withheld and only the approximate timing of the storms is supplied, we can estimate the top node saturation to within 3.4% in saturation terms, compared to 1.4% under the ideal conditions of the reference setup and 19% for the prior solution.

Note that the rms prior error increases steadily over the 2-week assimilation interval when precipitation measurements are withheld. This is not surprising since, in the absence of rainfall inputs, the prior solution is governed by a single, 2-week long drydown. For longer assimilation intervals the

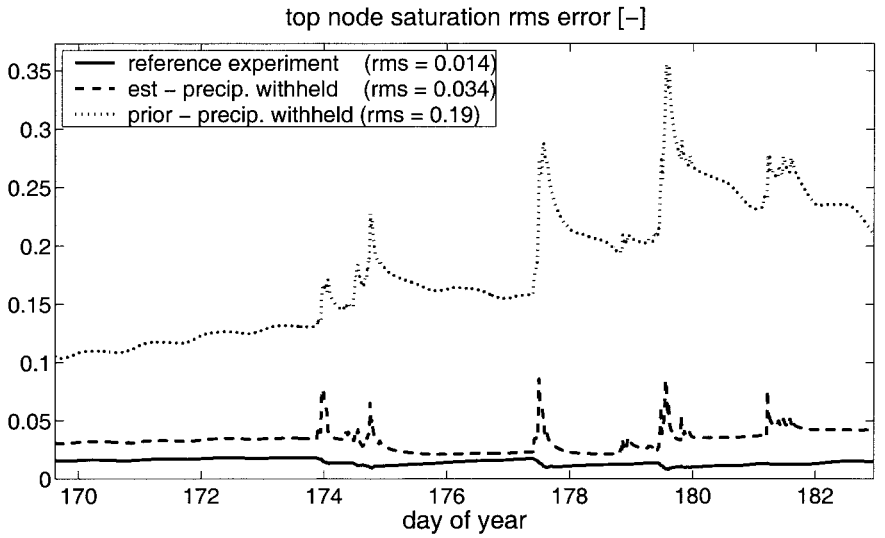

Figure 5. Areally averaged root-mean-square (rms) errors when precipitation is withheld. The solid line shows the rms errors from the reference experiment in which the precipitation data are provided to the assimilation algorithm. The dashed line shows the rms errors when the precipitation data are withheld. The dotted line shows the rms errors for the corresponding prior solution. The legend provides the temporal mean of the rms errors.

prior saturation error reaches a plateau since the saturation error is bounded.

Also note that the rms error in the top node saturation increases around precipitation events, reflecting the fact that the detailed temporal structure of these events cannot be resolved from brightness data that are available only once daily. To illustrate this point, we plot in Figure 6 the observed precipitation and the corresponding model error estimates for a typical pixel during a precipitation event. In this case the model error estimate is effectively an estimate of precipitation.

Although daily radio brightness measurements contain some information about relative storm volume, it is clear that they do not tell us much about the detailed temporal structure of individual storms. The model error estimates are smoother and generally lower in magnitude than the observed precipitation. This reflects the fact that each brightness observation only provides information on the cumulative effect of precipitation that occurred since the last measurement. The shape of the estimated precipitation events (i.e., estimated soil moisture model errors) shown in Figure 6 reflects the temporal correlation structure imposed by the model error statistics rather than information contained in the brightness measurements. This implies that fine-scale (temporal) rainfall intensity is unidentifiable from brightness measurements. Nevertheless, the soil moisture, which is the quantity of primary interest, can be accurately estimated from brightness data alone.

5.3. Multiple Assimilation Intervals

Perhaps the greatest advantage of variational data assimilation methods over Kalman filters is their ability to derive op-

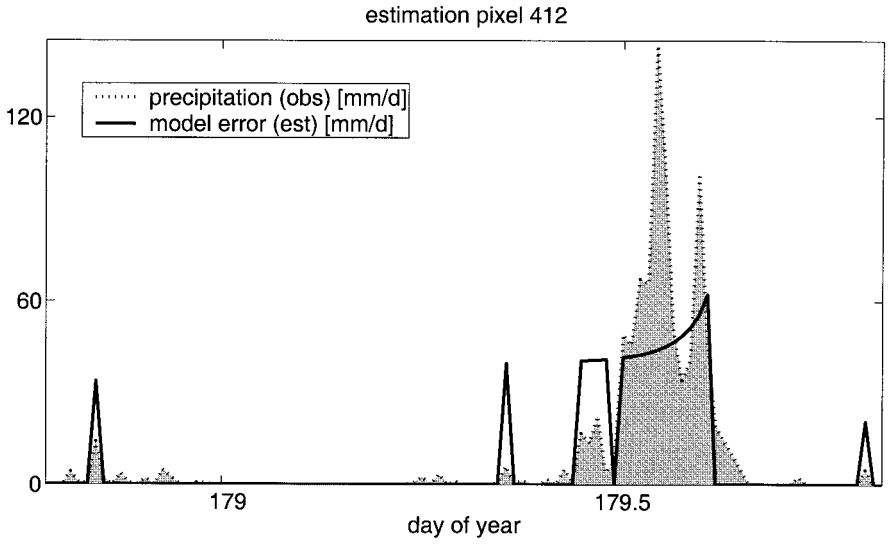

Figure 6. Model error estimates when precipitation data are withheld, plotted for a typical estimation pixel during the last major precipitation event.

timal estimates without explicitly computing an estimation error covariance matrix. However, in an operational setup, this can also be a disadvantage. If we want to use the variational method with a continual stream of new data, we have to block the data into a series of consecutive assimilation intervals $I_n \equiv [t_{n-1}, t_n]$ of finite length. To initialize the assimilation interval I_n, we must supply the prior (or first guess) mean value and covariance of the initial condition such that information gained in the past is reflected. The mean of the initial condition for I_n is readily set equal to the estimate at the end of the preceding interval I_{n-1}. Likewise, its covariance should equal the corresponding estimation error covariance at the end of the preceding interval I_{n-1}. However, this covariance is too expensive to compute in practice which forces us to adopt approximations.

In this section we study the impact of such approximations in a series of three synthetic experiments which focus on the issues of assimilation interval length and initialization. These experiments have the same specifications as the reference experiment except that the 2-week study period is divided into multiple assimilation intervals. One option to initialize interval I_n is to use the initial condition covariance that we specified in the single interval reference experiment. However, this single interval covariance will generally be unrealistically large when used to initialize I_n since it does not account for all the information obtained from data collected through I_{n-1}. It is better to make the initial condition covariance of I_n a scaled fraction of the single interval covariance. In our synthetic experiments the covariance scaling factor used for a given interval depends on the interval's sequential order and duration. Covariances for earlier intervals are higher because they rely on smaller amounts of historical data. This scaling does not change the associated correlation lengths.

In particular, in our first experiment (A) the 2-week study period is divided into three assimilation intervals containing 5, 3, and 4 radio brightness observation times, respectively (Figure 4). These intervals correspond roughly to interstorm periods. The corresponding initial condition covariances are 0.63, 0.19, and 0.17 of the single interval value. In our second and third experiments (B and C, respectively) we divide the study period into 12 short assimilation intervals, each containing only a single radio brightness measurement (Figure 4). In Experiment B the observation times are at the end of each interval, while in Experiment C they are at the beginning of each interval. In experiments B and C the initial condition covariance is set equal to $0.5n^{-1}$, where $n = 1, \ldots, 12$ indicates the sequential order of the interval.

Plate 2 shows the rms errors of the estimated top node saturation for the three assimilation interval experiments, together with the rms error obtained from the reference experiment. For Experiment A the time-averaged top node saturation rms error is 3%, and for Experiments B and C the corresponding time-averaged rms errors are 3.2 and 3.8%, respectively. These numbers compare to a time-averaged rms error of 2.9% in the reference experiment. As might be expected, the rms error increases as the number of assimilation intervals increases. With each additional assimilation interval we introduce more approximations by decoupling the data subsets associated with separate intervals and by naively reinitializing the initial condition covariance. This reduces the algorithm's ability to optimally combine the information contained in different intervals.

Perhaps the most interesting aspect of our assimilation interval study is the difference in Experiments B and C. In Experiment B, errors in specifying the initial condition covariance are introduced nearly 24 hours before the observation time. During the intervening period, time-dependent micrometeorological forcings gradually overwhelm the influence of the erroneous initial condition covariance. As a result, the comparison between predicted and measured radio brightness values, which occurs at the observation time, is relatively unaffected by errors in the initial condition covariance. By contrast, in Experiment C, poorly specified initial condition covariances have a greater influence on the update, which occurs only a few hours into the assimilation interval. Note that Experiment C corresponds closely to a Statistical Interpolation scheme, which updates estimates with measurements at the same time as the approximate covariance is specified independent of the dynamics of the system [*Daley*, 1991].

Overall, the estimates for the reference experiment and for Experiments A and B are very similar. This suggests that the suboptimality introduced by the naive reinitialization is not severe, provided the assimilation interval configuration allows errors in the initial condition covariance to evolve for at least one day. Given this, it is important to note that the algorithm's computational requirements decrease substantially if we use multiple but shorter assimilation intervals. For the case considered here division of the study period into three assimilation intervals decreases the computer time by nearly two thirds. The savings are even more dramatic if we use 12 assimilation intervals. This decrease in computer time reflects the fact that initial condition updates are less computationally demanding than model error updates.

Although the multiple interval covariance reinitialization technique used here is naive and the estimation process is not strictly optimal, the estimates are quite close to the optimal estimates of the reference experiment. For a small sacrifice in optimality, the savings in computational effort are sufficient to make assimilation intervals of a few days an attractive option.

6. Summary and Conclusions

In this work we have developed and investigated an optimal four-dimensional data assimilation algorithm for hydrologic applications. This algorithm is based on a land surface model that describes changes in soil moisture, soil temperature, and canopy temperature over time and three spatial dimensions. The land model captures key physical processes at the land-atmosphere boundary while meeting the stringent computational requirements of a practical data assimilation system. Computational efficiency has been achieved by dividing the model domain into laterally uncoupled one-dimensional vertical columns. These columns are related through horizontal correlations in the various inputs which force the model. The use of a column-oriented model also makes it possible to estimate profiles of soil moisture and related variables.

Estimates of the land surface states are derived directly from L-band passive microwave measurements, which are related to the states by a radiative transfer model. Other types of remotely sensed observations could be included with suitable modifications. The estimation process is formulated as a variational optimization problem which accounts for both model and measurement errors. The solution is obtained with the

iterated indirect representer algorithm technique, which is applied here for the first time to a hydrologic problem.

We have used our data assimilation algorithm to investigate several design issues which are important in practical applications. These issues deal with downscaling of radio brightness observations, the role of precipitation measurements in soil moisture estimation, and the configuration of assimilation intervals. Our synthetic downscaling experiments show that it is feasible to estimate land surface variables such as soil moisture at horizontal resolutions smaller than the resolution of available radio brightness observations. The additional information gained in this case comes from micrometeorological, land cover, and soil texture data provided at smaller spatial scales. Even for downscaling ratios of one to sixteen (each observation pixel contains sixteen estimation pixels), it is possible to capture many of the fine-scale features of the true fields. This implies that brightness images with a resolution of 40–50 km may be sufficient to support soil moisture estimates on a scale of 10 km, provided that sufficiently accurate model inputs are available at the finer scale.

In other synthetic experiments we have demonstrated the ability of the assimilation algorithm to satisfactorily estimate soil moisture even if quantitative precipitation data are not available. This is possible because radio brightness data are sensitive to precipitation-induced changes in soil moisture. Although soil moisture estimates improve if both radio brightness and precipitation data are available, it may be sufficient in practice to provide the assimilation algorithm with information on the timing of rainfall events.

In our final set of synthetic experiments we have assessed the influence of the length of the assimilation interval, and we have considered the operational problem of specifying an approximate initial condition covariance at the beginning of each interval. Our investigation of this issue indicates that it is best to match the assimilation interval approximately to an interstorm period. Moreover, each interval should be chosen to ensure that there is sufficient time for errors in the approximate initial condition covariance to dampen out before the next measurement is processed. If we observe these guidelines we can obtain nearly optimal estimates, even though computational limitations force us to naively reinitialize the assimilation intervals.

The next step in the development of our data assimilation approach is to conduct a field test, using ESTAR radio brightness observations collected during the SGP97 experiment [*Jackson et al*., 1999]. This will force us to reexamine issues relating to model error, both forcing errors and more pervasive errors in model structure. In future applications computation of selected estimation error variances should be included. This is within the capabilities of the variational approach, although it is a computationally demanding task. Estimation error variances will provide valuable insights about the accuracy of the estimates and can provide a basis for tests of hypotheses about error statistics and other algorithm inputs [*Reichle et al.*, 2001]. Data compression may help to lower the cost of computing estimation error variances, especially if the size of the estimation problem can be substantially reduced in this way.

In practice, the feasibility of retrieving subsurface moisture profiles from surface measurements depends on the accuracy and the physical realism of the land model and the associated error statistics. Since large-scale soil moisture profiles cannot be directly observed they can only be estimated by using the hydrologic model to propagate information downward from the surface. However, the time required for reliably estimating profiles in this way is significantly longer than the 2-week period considered here. One approach to decrease the required time could be to use physically consistent similarity profiles and estimate the scale parameters. Other problems with profile estimation are the scarcity of verification data and the fact that the connection between profile measurements at point locations and large-scale estimates from data assimilation remains unclear. Note, however, that large-scale profile estimates can be useful independent of their relationship to in situ measurements by providing physically consistent large-scale surface flux estimates which can be verified more easily. These issues clearly deserve more attention in future research.

It is likely that land surface estimates produced by our assimilation algorithm could be improved by adding other types of measurements, such as infrared observations of soil temperature and low-frequency microwave measurements (for example, C-band, 5.3 GHz). Multi-platform and multifrequency measurements can be naturally included in the assimilation algorithm if the measurement equation is modified accordingly.

The algorithm in its current implementation is not yet sufficiently robust for widespread practical applications. Convergence is difficult to achieve for large uncertainties in the initial condition and long assimilation intervals. It is not yet clear whether the difficulties arise because of the nature of land processes or because of possible deficiencies in our implementation, particularly of the model adjoint. More research is necessary to trace the source of these problems. Nevertheless, our results on downscaling and the role of precipitation for soil moisture estimation are valid independent of the choice of assimilation algorithm.

Finally, it should be noted that further improvements in computational efficiency may be required before continental-scale applications of hydrologic data assimilation become feasible. Given the highly nonlinear structure of land-atmosphere processes and the high complexity of real world applications, any large-scale land assimilation algorithm will inevitably be a compromise between realistic physical representations and computational feasibility. We believe that the approach outlined here strikes a reasonable balance on this continuum.

Appendix A: Downscaling Example

Both the state vector $Y(t)$ and its adjoint $\lambda(t)$ are defined on the scale of the estimation pixels and are related through the Euler-Lagrange equations. Consequently, we can illustrate how downscaling works by considering how the estimation-scale adjoint variable can be derived from observation-scale brightness data. Suppose the domain consists of four estimation pixels and the state vector is only given by the soil saturation, that is $Y = [S_1 \ldots S_4]^T$. Let us further assume that there is only one scalar brightness measurement $Z = T_B^{\text{obs}}$ at time t_1 which covers all four estimation pixels (Figure 1) and which has a measurement error variance σ_v^2.

By using the radiative transfer model of section 3, we can derive model predictions of the brightness temperature $T_{Bi}(S_i)$ on the scale of the estimation pixels. In order to compute the data misfit term $Z - M[Y]$, we need model predictions of brightness temperature on the coarser observation scale. This is easily achieved by averaging the fine-scale brightness predictions over the observation pixel. For L-band observations the arithmetic average is appropriate [*Drusch et*

al., 1999; *Liou et al*., 1998]. In our example the measurement operator is therefore $M[Y] = \frac{1}{4}\sum_{j=1}^{4} T_{Bj}$. Using the above in (5) yields for the ith estimation pixel ($i = 1, \ldots, 4$)

$$-\frac{\partial \lambda_i}{\partial t} = \sum_{j=1}^{4} \left.\frac{\partial \varphi_j}{\partial Y_i}\right|_{\hat{Y}} \lambda_j + \frac{1}{4}\left.\frac{\partial T_{Bi}}{\partial S_i}\right|_{\hat{Y}} \cdot \delta(t - t_1)\sigma_v^{-2}\left(T_B^{\text{obs}} - \frac{1}{4}\sum_{j=1}^{4} T_{Bj}\right). \quad (7)$$

Equation (7) shows how the algorithm distributes the information from the coarse-scale brightness observation to the finer scale of the adjoint (or equivalently the state). The observation-scale data misfit term $T_B^{\text{obs}} - \frac{1}{4}\sum_{j=1}^{4} T_{Bj}$ is weighted differently for each of the four estimation pixels. The respective weights depend on estimation pixel-scale micrometeorological, soil, and land cover data that affect the partial derivatives $\partial\varphi_j/\partial Y_i$ and $\partial T_{Bi}/\partial S_i$. See Figure 2 for an example of how different soil textures affect the microwave emissivity and ultimately the brightness temperature.

The brightness measurement T_B^{obs} enters (7) only at the observation scale. No off-line disaggregation is necessary. On the contrary, the fine-scale model predictions of brightness are aggregated to the observation scale in order to calculate the data misfit term. An estimate of the downscaled model error is obtained by substituting the downscaled adjoint variable into (6). The downscaled state estimate is then obtained by substituting the downscaled model error into (4).

Acknowledgments. This research was sponsored by the NASA Land Surface Hydrology Program (NRA-98-OES-11). We would also like to thank the Earth System Science Center at Pennsylvania State University and the Oklahoma Mesonet for their invaluable data. Alan Willsky, John Galantowicz, and two anonymous reviewers provided many helpful comments and suggestions.

References

Bennett, A. F., *Inverse Methods in Physical Oceanography*, Cambridge Univ. Press, New York, 1992.

Bennett, A. F., Inverse methods and data assimilation, in Lecture Notes of the 1999 Summer School at the College of Oceanic and Atmospheric Sciences at Oregon State University, Corvallis, Oreg., 1999.

Bennett, A. F., B. S. Chua, D. E. Harrison, and M. J. McPhaden, Generalized inversion of tropical atmosphere-ocean data and a coupled model of the tropical Pacific, *J. Clim*., *11*(7), 1768–1792, 1998.

Calvet, J.-C., J. Noilhan, and P. Bessemoulin, Retrieving the root-zone soil moisture from surface soil moisture or temperature estimates: A feasibility study based on field measurements, *J. Appl. Meteorol*., *37*(4), 371–386, 1998.

Castelli, F., D. Entekhabi, and E. Caporali, Estimation of surface heat flux and an index of soil moisture using adjoint-state surface energy balance, *Water Resour. Res*., *35*(10), 3115–3125, 1999.

Clapp, R. B., and G. M. Hornberger, Empirical equations for some soil hydraulic properties, *Water Resour. Res*., *14*, 601–604, 1978.

Courtier, P., J. Derber, R. Errico, J.-F. Louis, and T. Vulicevic, Important literature on the use of adjoint, variational methods and the Kalman filter in meteorology, *Tellus Ser. A*, *45*, 342–357, 1993.

Daley, R., *Atmospheric Data Analysis*, Cambridge Univ. Press, New York, 1991.

Drusch, M., E. F. Wood, and C. Simmer, Up-scaling effects in passive microwave remote sensing: ESTAR 1.4 GHz measurements during SGP '97, *Geophys. Res. Lett*., *26*(7), 879, 1999.

Entekhabi, D., H. Nakamura, and E. G. Njoku, Solving the inverse problem for soil moisture and temperature profiles by sequential assimilation of multifrequency remotely sensed observations, *IEEE Trans. Geosci. Remote Sens*., *32*(2), 438–448, 1994.

Galantowicz, J. F., D. Entekhabi, and E. G. Njoku, Tests of sequential data assimilation for retrieving profile soil moisture and temperature from observed L-band radio brightness, *IEEE Trans. Geosci. Remote Sens*., *37*(4), 1860, 1999.

Gelb, A., (Ed.), *Applied Optimal Estimation*, MIT Press, Cambridge, Mass., 1974.

Houser, P. R., W. J. Shuttleworth, J. S. Famiglietti, H. V. Gupta, K. H. Syed, and D. C. Goodrich, Integration of soil moisture remote sensing and hydrologic modeling using data assimilation, *Water Resour. Res*., *34*(12), 3405–3420, 1998.

Hu, Z., and S. Islam, Prediction of ground surface temperature and soil moisture content by the force-restore method, *Water Resour. Res*., *31*, 2531–2539, 1995.

Jackson, T. J., D. M. Le Vine, A. Y. Hsu, A. Oldak, P. J. Starks, C. T. Swift, J. D. Isham, and M. Haken, Soil moisture mapping at regional scales using microwave radiometry: The Southern Great Plains Hydrology Experiment, *IEEE Trans. Geosci. Remote Sens*., *37*(5), 2136–2151, 1999.

Jackson, T. J., P. E. O'Neill, and C. T. Swift, Passive microwave observation of diurnal surface soil moisture, *IEEE Trans. Geosci. Remote Sens*., *35*(5), 1210–1222, 1997.

Katul, G. G., O. Wendroth, M. B. Parlange, C. E. Puente, M. V. Folegatti, and D. R. Nielsen, Estimation of in situ hydraulic conductivity function from nonlinear filtering theory, *Water Resour. Res*., *29*(4), 1063–1070, 1993.

Liou, Y. A., E. J. Kim, and A. W. England, Radiobrightness of prairie soil and grassland during dry-down simulations, *Radio Sci*., *33*(2), 259–265, 1998.

McLaughlin, D., Recent advances in hydrologic data assimilation, *U.S. Natl. Rep. Int. Union Geod. Geophys. 1991–1994*, *Rev. Geophys*., *33*, 977–984, 1995.

Parlange, M. B., G. G. Katul, M. V. Folegatti, and D. R. Nielsen, Evaporation and the field scale soil water diffusivity function, *Water Resour. Res*., *29*(4), 1279–1286, 1993.

Reichle, R. H., Variational Assimilation of Remote Sensing Data for Land Surface Hydrologic Applications, Ph.D. dissertation, Dep. of Civ. and Environ. Eng., Mass. Inst. of Technol., Cambridge, Mass., 2000.

Reichle, R. H., D. B. McLaughlin, and D. Entekhabi, Variational data assimilation of microwave radio brightness observations for land surface hydrology applications, *IEEE Trans. Geosci. Remote Sens.*, in press, 2001.

Thépaut, J.-N., and P. Courtier, Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model, *Q. J. R. Meteorol. Soc*., *117*, 1225–1254, 1991.

Ulaby, F. T., R. K. Moore, and A. K. Fung, *Microwave Remote Sensing*, vol. I–III, Artech House, Norwood, Mass., 1986.

Xue, Y., P. J. Sellers, J. L. Kinter, and J. Shukla, A simplified biosphere model for global climate studies, *J. Clim*., *4*, 345–364, 1991.

D. Entekhabi and D. B. McLaughlin, Ralph M. Parsons Laboratory, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

R. H. Reichle, NASA Goddard Space Flight Center, Code 974, Greenbelt, MD 20771, USA. (reichle@alum.mit.edu)

(Received March 21, 2000; revised November 26, 2000; accepted April 25, 2001.)